arXiv:1611.07161v1 [stat.ML] 22 Nov 2016

OPTIMAL LEARNING FOR STOCHASTIC OPTIMIZATION WITH NONLINEAR PARAMETRIC BELIEF MODELS

XINYU HE* AND WARREN B. POWELL†

Abstract. We consider the problem of estimating the expected value of information (the knowledge gradient) for Bayesian learning problems where the belief model is nonlinear in the parameters. Our goal is to maximize some metric, while simultaneously learning the unknown parameters of the nonlinear belief model, by guiding a sequential experimentation process which is expensive. We overcome the problem of computing the expected value of an experiment, which is computationally intractable, by using a sampled approximation, which helps to guide experiments but does not provide an accurate estimate of the unknown parameters. We then introduce a resampling process which allows the sampled model to adapt to new information, exploiting past experiments. We show theoretically that the method converges asymptotically to the true parameters, while simultaneously maximizing our metric. We show empirically that the process exhibits rapid convergence, yielding good results with a very small number of experiments.

Key words. Optimal Learning, Knowledge Gradient, Nonlinear Parametric Model

AMS subject classifications. 62L05, 62F07, 62F12, 62F15

1. Introduction.

We consider the following problem: we have a function over a finite number of alternatives, where the expression of the function is known except for some unknown parameters. Our goal is both to learn the correct parameters and to find the alternative that maximizes the function. Information can be collected by measuring the function value at any chosen alternative. However, the measurements are noisy and expensive, and we only have a limited budget to evaluate these alternatives. After exhausting the budget, we have to provide the best estimate of the unknown parameters and identify the best alternative to maximize our metric.

The optimization of functions with noisy measurements has been broadly studied in the stochastic search community. [27] provides an extensive review of different methods for stochastic search. Some commonly used algorithms include gradient estimation [12], response surface methodology [3], and metaheuristics such as tabu search and genetic algorithms ([24], also see [4] for more on metaheuristics). These methods are usually designed for functions that are relatively easy to evaluate, while our focus is on problems where a function evaluation is expensive (and we do not have access to derivatives).

One of the main challenges of the problems with expensive experiments is known as "*exploration vs exploitation*" (see Chapter 12 of [25]), which requires a balance of exploiting the current optimal solutions and exploring with uncertain outcomes to learn the problem. Specifically, for a finite number of alternatives, some heuristics to strike this balance include epsilon-greedy ([29], also see [26] for convergence analysis), interval estimation [21], and Chernoff interval estimation [28]. These methods usually perform well on small problems but have scaling issues in higher dimensions. More sophisticated methods to collect information efficiently include Gittins indices [13], upper confidence bound methods [2], and expected improvement ([20], also see [14] for recent work).

*Department of Electrical Engineering, Princeton University, Princeton, NJ, 08544 (xinyuhe@princeton.edu).

†Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ, 08544 (powell@princeton.edu).

This problem is an extension of the ranking and selection (R&S) problem (see [22, 18] and the references cited there), which focuses on finding the optimal alternative, as opposed to learning the values of an underlying, parametric belief model. There have been two major approaches to R&S problems: the frequentist approach and the Bayesian approach. The frequentist approach assumes that information only comes from observations and uses the distribution of the observed data to estimate the parameters (see, e.g., [16, 22, 1]). The Bayesian approach begins with a probabilistic belief model about the true parameters which are updated as new observations are collected (see, e.g., [5, 8]). For the Bayesian approach, there are mainly two directions: the Optimal Computing Budget Allocation (OCBA) (see, e.g., [6, 5, 17]), which tries to maximize the posterior probability of correct selection, and the Value of Information (VoI) procedures (see, e.g., [9, 10]), which maximizes the improvement in a single experiment. Below we give a brief review of the VoI procedures, with special attention given to the knowledge gradient [10].

Since every experiment is expensive or time consuming (consider for example, the laboratory experiments in materials science that might take days or weeks), it is important to maximize the value of each measurement. [15] proposes the idea of computing the marginal value of information from a single experiment. Based on this, [10] extends the idea using the Bayesian approach, and presents the Knowledge Gradient (KG) policy. It analyzes the case where all alternatives are independent, and also proves that KG is the only stationary policy that is both myopically and asymptotically optimal. [11] adapts the knowledge gradient to handle correlations in the beliefs about discrete alternatives. Both [10] and [11] use lookup table belief models, and become computationally expensive when the number of alternatives becomes large, as typically happens when an experimental choice involves multiple dimensions. [23] is the first paper that studied KG using a parametric belief model. It assumes that the belief model is linear in some unknown parameters, and imposes the uncertainty of the function values onto the parameters. This strategy reduces the number of parameters to be estimated from the number of alternatives (required by a lookup table belief model) to a much lower-dimensional parameter vector. While the knowledge gradient is easy to compute for belief models that are linear in the parameters, nonlinear belief models are computationally intractable. [7] studies the more general nonlinear parametric model, and proposes the Knowledge Gradient with Discrete Priors (KGDP), which overcomes the computational issues, but requires that we represent the uncertainty about the unknown parameters using a finite number of samples, where one of the samples has to be the true parameter. KGDP is able to handle any nonlinear belief model, but the assumption that one of the candidates is correct (we refer to this assumption later as the *truth from prior* assumption) seems too strong in most real world settings, especially when the parameters have four or more dimensions. Moreover, [7] also fails to give any theoretical proof of convergence of KGDP given the truth from prior assumption.

In this paper, we present a resampling algorithm that works with KGDP to find not only the best alternative but also the correct parameters without the assumption of one candidate being the truth. We start with several potential candidates, and use the information from the measurements to guide resampling and discover more probable candidates. Regarding the objective of finding the correct parameter, we propose a new metric for calculating the knowledge gradient that minimizes the entropy of the candidate parameters. We also prove the convergence of both the non-resampling (with truth from prior) and the resampling KGDP algorithms. In the experimental

section, we apply our algorithm to a real problem in materials science: optimizing the kinetic stability of an experimental problem involving the control of a water-oil-water (W/O/W) nano emulsion [7]. Compared with [7], our algorithm shows significant improvements.

In order to find the correct parameter, we need to deal with the parameter space (denoted as θ space) besides the alternative space (denoted as x space). Here arises a dual problem: in x space, we solve a maximization problem to figure out which alternative maximizes some performance metric (strength, conductivity, reflexivity); in θ space, we solve another optimization problem to locate the most promising candidates. The second optimization problem is solved via minimizing the mean square error (MSE).

To the best of our knowledge, this is the first optimal learning paper that addresses the dual problems of objective maximization with parameter identification for problems with parametric belief models (other literatures with similar dual-objective formulations usually assume that experiments are inexpensive, e.g, [19]). Previous papers only concentrate on discovering the optimal alternative, but in many real world situations, scientists also care about getting accurate estimates of the parameters. For example, materials scientists may want to know both the tunable variables of an experiment (such as the temperature, the pressure, and the density), and some unknown and uncontrollable parameters, which can help them understand the intrinsic nature of the physical phenomena. Our work in this paper provides a powerful solution to this issue.

This paper makes the following major contributions:

- We present a resampling algorithm in the parameter space, which works with the optimal learning methods in the alternative space, to discover both the optimal alternative and the correct parameters.
- We propose a new metric to calculate the knowledge gradient, which focuses on reducing the uncertainty of parameters by maximizing entropic loss.
- We prove the asymptotic optimality of both the nonresampling (with truth from prior) and the resampling algorithms using either traditional function oriented metric or our new entropy oriented metric.
- We show experimentally that our resampling algorithm has impressive performance in both finding the optimal alternative and estimating the parameters.

The paper is organized as follows. Section 2 reviews the principles of optimal learning based on maximizing the value of information, along with the knowledge gradient with discrete priors (KGDP) proposed in [7]. In Section 3, we introduce the resampling algorithm as well as the entropy oriented metric for calculating the knowledge gradient. Section 4 provides the asymptotic convergence proof of KGDP without resampling under truth from prior assumption, and Section 5 proves asymptotic convergence of KGDP with resampling. We show empirical results on both a simulated and a real-world problem in Section 6, and finally conclude in Section 7.

2. Knowledge Gradient for Nonlinear Belief Models.

In this section, we first review the ranking and selection problem along with the knowledge gradient policy, and then introduce the model of Knowledge Gradient with Discrete Priors.

2.1. Ranking and Selection (R&S) Problem.

Suppose we have a finite number of alternatives $\mathcal{X} = \{x_1, x_2, ..., x_M\}$. Each alternative $x \in \mathcal{X}$ is associated with a true utility μ_x, which is presumed unknown to us. The goal is to determine the x that has the largest utility through a budget of N sequential measurements. At time n (i.e, after n measurements), suppose we decide to query alternative x^n according to some

policy, and our measurement is $\hat{y}^{n+1}$, where the superscripts imply that $\hat{y}^{n+1}$ will be unknown until the $(n+1)$-th measurement. We assume the inherent noise (W^n) in each measurement follows a normal distribution with zero mean and variance σ^2, where σ is known to us. That is, for $n = 0, 1, ..., N-1$,

$$\hat{y}^{n+1} = \mu_{x^n} + W^{n+1},$$

where $W^{n+1} \sim \mathcal{N}(0, \sigma^2)$.

For each $x \in \mathcal{X}$, we use θ_x^n as our estimate of the true utility μ_x after n experiments. Our goal is to select the x with the highest estimate after N measurements, i.e.,

$$x^N = \underset{x \in \mathcal{X}}{\operatorname{argmax}} \theta_x^N.$$

2.2. Knowledge Gradient. Let S^n denote the state of knowledge at time n, and let $V^n(S^n)$ be the value of information if we are in state S^n. In the R&S problem, we have $V^n(S^n) = \max_{x' \in \mathcal{X}} \theta_{x'}^n$. The transition from S^n to S^{n+1} occurs when we take a measurement at $x^n = x$ and observe $\hat{y}^{n+1}$, with

$$V^{n+1}(S^{n+1}(x)) = \max_{x' \in \mathcal{X}} \theta_{x'}^{n+1}(x).$$

At time n, $\theta_{x'}^{n+1}(x)$ is a random variable since it depends on the noise W^{n+1}. We would like to maximize our expected incremental information from the next measurement, which we call the *knowledge gradient*. At time n, the knowledge gradient of x is defined as:

$$\begin{aligned} \nu^{KG,n}(x) &= \mathbb{E}^n \left[V^{n+1}(S^{n+1}(x)) - V^n(S^n) \right] \\ &= \mathbb{E}^n \left[\max_{x' \in \mathcal{X}} \theta_{x'}^{n+1}(x) | S^n = s, x^n = x \right] - \max_{x' \in \mathcal{X}} \theta_{x'}^n. \end{aligned} \tag{2.1}$$

In each iteration, the *Knowledge Gradient* policy measures the alternative with the largest knowledge gradient:

$$x^{KG,n} = \underset{x \in \mathcal{X}}{\operatorname{argmax}} \nu^{KG,n}(x).$$

This is a myopic policy that maximizes the value of information from a single experiment [10].

The knowledge gradient Equation (2.1) is easy to calculate when using a lookup table belief model ([10] and [11]), or when the belief model is linear in the parameters [23], with the form $f(x;\theta) = \theta_0 + \theta_1\phi_1(x) + ... + \theta_n\phi_n(x)$. We run into problems, however, if the function is nonlinear in the parameter vector θ, since we have to incorporate the updating mechanism in the representation of $\theta_{x'}^{n+1}(x)$ in Equation (2.1). When this is imbedded within both the maximum and expectation operator, the resulting expression becomes computationally intractable. We address this issue in the next section.

2.3. Knowledge Gradient with Discrete Priors (KGDP). The Knowledge Gradient with Discrete Priors (KGDP) [7] is designed to handle parametric belief models that are nonlinear in the parameters. Suppose we have a function $f(x;\theta)$, where x is the alternative and θ represents the unknown parameters. The expression of $f(x;\theta)$ is known to us except for the unknown θ. Let θ^* denote the true parameter. Our goal is to find the x that maximizes the true function $f(x;\theta^*)$. KGDP assumes that $f(x;\theta^*)$ can be approximated as a convex combination of L candidates:

$$f(x;\theta^*) \approx \bar{f}^n(x) = \sum_{i=1}^{L} f(x;\theta_i) p_i^n,$$

where θ_i's are known as the candidates, with p_i^n's as their probabilities at time n.

KGDP requires that one of the candidates is equal to or very close to the truth. We call this the *truth from prior* assumption, while its opposite, *truth not from prior*, recognizes that the true θ may not belong to our sampled set. As we take measurements, the θ's are fixed, while the probabilities p_i's get updated after each iteration, which means that the belief is naturally conjugate. The state variable S^n is defined as the probability vector, $S^n = (p_1^n, p_2^n, ..., p_L^n,)$. The KGDP score of each alternative x can be calculated according to Equation (2.1) (we review this formula in Section 3.5.1). We calculate the KGDP scores of all alternatives, and select the one with the highest score to evaluate next. After we exhaust the budget, the true parameters can be estimated by choosing the most probable candidate, i.e, $\hat{\theta}^* = \theta_i$, where $i = \text{argmax}_{l \in [L]} \, p_l^N$.

KGDP provides a powerful approach to handling nonlinear parametric models if we use a sampled belief model for the unknown parameters. However, the complexity of KGDP grows quickly as the number of candidates increases. In the experiments described in [7], it can handle at most tens of candidates. Moreover, the truth from prior assumption may be a reasonable approximation if θ has one or two dimensions, but in higher dimensions it is unlikely that this would be the case. This is particularly problematic when we are interested in finding the best possible estimate of the parameters themselves. For this reason, we propose a resampling algorithm that can adaptively find more promising parameter candidates.

3. Optimal Learning for Parameters and Alternatives. Our dual-objective optimization problem requires handling the alternative space and the parameter space simultaneously. Note that while finding the best alternative and learning the best estimate of the parameter vector each have their own metrics, they are not conflicting goals, since the best estimate of the parameters can lead to the best design.

Traditionally, the optimal learning literature has focused on optimizing some metric such as choosing the best parameters to maximize the strength of a material or minimize the number of defects. When we use parametric belief models, we can achieve these goals by finding good estimates of the unknown parameters, and then optimizing a (deterministic) nonlinear model based on the estimates. To capture the value of learning the unknown parameters correctly, we can replace our original metric with the entropy of the belief vector.

3.1. The Resampling Scheme. As in KGDP, we still keep a set of L candidates, but they change over time. Let $\mathbb{L}^n = \{\theta_1^n, ..., \theta_L^n\}$ denote the candidate set at time n, and $\vec{p^n} = \{p_1^n, ..., p_L^n\}$ denote the probabilities. Assume that the true function $f(x; \theta^*)$ can be approximated as

$$f(x; \theta^*) \approx \bar{f}^n(x) = \sum_{i=1}^{L} f(x; \theta_i^n) p_i^n.$$

The candidates are no longer fixed and resampling happens periodically, in which the least probable candidates get replaced by more probable ones according to the whole history of experiments. Hence, we need to search two spaces: in the outer loop, we work in the x space to collect as much information as possible; in the inner loop, we search the θ space for more promising θ's.

We use the following process for testing our resampling strategy. We start by choosing a large sample of K (say $K = 1,000,000$) realizations of θ's, and then choose one of these as the truth. We call the K samples the *large pool*, denoted as

$\mathbb{K}$. Then, we choose a small sample set $\mathbb{L}$ as our sampled prior. The idea is that with such a large base population, it will be very unlikely that one of the θ's in our sampled prior will be the truth (or even close to it).

As the candidates' probabilities change over time, under certain conditions resampling is triggered. We use mean square error (MSE) as the criterion to choose θ's in the resampling steps.

3.2. Mean Square Error (MSE) Formulation. We first introduce the formulation of the MSE problem, which is equivalent to the maximum likelihood estimation (MLE) in Gaussion noise settings. Suppose that after n measurements, the history of experimental results is $\mathcal{F}^n = \sigma\{(x^0, \hat{y}^1), ..., (x^{n-1}, \hat{y}^n)\}$. The likelihood of each θ_i is given by:

$$\mathcal{L}(\theta_i|\mathcal{F}^n) = \prod_{j=1}^{n} \exp\left(-\frac{[\hat{y}^j - f(x^{j-1};\theta_i)]^2}{2\sigma^2}\right).$$

By taking the logarithm of $\mathcal{L}(\theta_i|\mathcal{F}^n)$, we have the MSE formula for θ_i as

$$MSE(\theta_i|\mathcal{F}^n) = \frac{1}{n}\sum_{j=1}^{n}[\hat{y}^j - f(x^{j-1};\theta_i)]^2.$$

The idea is that each time resampling is triggered, we calculate the MSE of all θ's in the large pool $\mathbb{K}$. Although K can be arbitrarily large, this calculation is considered much faster and more efficient compared to the expensive physical experiments. Then resampling is conducted in the sub-level set of the MSE function $MSE(\theta)$. We define a threshold δ, and all θ's with MSE smaller than δ form the sub-level set $\mathbb{S}^n$:

$$\mathbb{S}^n = \{\theta \in \mathbb{K}|MSE(\theta|\mathcal{F}^n) \leq \delta\}.$$

Alternatively, we can define a *small pool* of a certain size R that contains the R θ's with the smallest MSE, namely

$$\mathbb{S}^n = \{\theta \in \mathbb{K}|MSE(\theta|\mathcal{F}^n) \leq MSE_R\}.$$

where MSE_R denotes the value of the R-th smallest MSE. In other words, the small pool $\mathbb{S}^n$ contains the most probable θ's up to now. We resample from the small pool a number of new candidates as needed.

On the one hand, this method rules out the majority of unlikely samples and avoids exploring too large a space inefficiently; on the other hand, it supplies a small pool containing sufficiently promising samples to resample from, thus providing plenty of exploration compared with simply selecting the L samples with the smallest MSE.

3.3. Updating the Probabilities of Candidates. For candidate θ_i^n, suppose its prior probability is p_i^n. After the $(n+1)$-th measurement given by $\hat{y}^{n+1} \sim \mathcal{N}(f(x^n;\theta_i^n), \sigma^2)$, the likelihood of $\hat{y}^{n+1}$ is given by

$$g_Y(\hat{y}^{n+1}|\theta^* = \theta_i^n) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{[\hat{y}^{n+1} - f(x^n;\theta_i^n)]^2}{2\sigma^2}\right).$$

By Bayes' rule, the posterior probability is proportional to the prior probability times the likelihood:

$$p_i^{n+1} \propto g_Y(\hat{y}^{n+1}|\theta^* = \theta_i^n)p_i^n = \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{[\hat{y}^{n+1} - f(x^n;\theta_i^n)]^2}{2\sigma^2}\right)p_i^n. \tag{3.1}$$

After normalization, the updating rule of p_i is given by

$$p_i^{n+1} = \frac{\exp\left(-\frac{[\hat{y}^{n+1} - f(x^n;\theta_i^n)]^2}{2\sigma^2}\right) p_i^n}{\sum_{l=1}^{L} \exp\left(-\frac{[\hat{y}^{n+1} - f(x^n;\theta_l^n)]^2}{2\sigma^2}\right) p_l^n}. \tag{3.2}$$

After the $(n+1)$-th measurement, if resampling is not triggered, p_i^{n+1} can be updated as Equation (3.2) for $i = 1, 2, ..., L$. Otherwise, we should calculate the likelihood of θ_i by all the previous measurements, given by

$$g_Y(\hat{y}^1, ..., \hat{y}^{n+1} | \theta^* = \theta_i^{n+1}) = \prod_{j=0}^{n} \exp\left(-\frac{[\hat{y}^{j+1} - f(x^j;\theta_i^{n+1})]^2}{2\sigma^2}\right). \tag{3.3}$$

Since initially all parameters are associated with equal probability, if resampling happens at time $(n+1)$, p_i^{n+1} is given by

$$p_i^{n+1} = \frac{\prod_{j=0}^{n} \exp\left(-\frac{[\hat{y}^{j+1} - f(x^j;\theta_i^{n+1})]^2}{2\sigma^2}\right)}{\sum_{l=1}^{L} \prod_{j=0}^{n} \exp\left(-\frac{[\hat{y}^{j+1} - f(x^j;\theta_l^{n+1})]^2}{2\sigma^2}\right)}, \tag{3.4}$$

where i and l index the set of candidates after resampling.

3.4. Resampling Procedure. Resampling is triggered under either of two conditions: 1) the same set of candidates have been used for n^{resamp} iterations; 2) more than $L/2$ candidates have probabilities lower than ϵ. The resampling process goes as follows.

1. When resampling is triggered, we first remove the candidates with low probabilities. Denote the set as $\mathbb{L}_{rm}^n$:

$$\mathbb{L}_{rm}^n = \{\theta \in \mathbb{L}^n | p^n(\theta) \leq \epsilon\}.$$

Note that if $p^n(\theta) > \epsilon$ for all $\theta \in \mathbb{L}^n$, we still select a small portion (say 1 or 2) of the least probable ones as $\mathbb{L}_{rm}^n$. In this way, we can avoid getting stuck in a set of unlikely candidates (for example, in an extreme case where we have L identical but wrong candidates, we get stuck if not dropping anyone).

2. Then we calculate the MSE of each θ in the large pool by Equation (3.1), and select R ones with the smallest MSE to form the small pool.

3. Next, calculate the likelihood of each θ in the small pool given by Equation (3.3). We then use the likelihoods as weights and do weighted sampling without replacement to select $|\mathbb{L}_{rm}^n|$ θ's and add them to the candidate set.

4. Once the candidates are updated, we update their probabilities accordingly using the whole measurement history according to Equation (3.4).

5. Finally, we check if the current set can still trigger the resampling conditions, since the new set may still contain over $L/2$ candidates with probabilities lower than ϵ. If not, resampling finishes. Otherwise, repeat the resampling process.

A detailed flowchart is given in Section 3.6.

3.5. Evaluation Metrics. We introduce two Knowledge Gradient related policies for choosing the next alternative to measure. Corresponding to our dual objectives, they focus on maximizing the performance metric and learning the parameter respectively. The first one, initially given in [7], is function value oriented and hence denoted as KGDP-f. The second one, denoted as KGDP-H, focuses on minimizing the entropy of the belief vector $\vec{p}^n = \{p_1^n, ..., p_L^n\}$, which leads to a better estimate of θ^*, from which we can then optimize the original function $f(x;\theta)$.

3.5.1. KGDP-f. We can measure the expected incremental function value as in [7], where the formula of the KGDP-f was originally given. At time n, we define S^n as the probability vector $(p_1^n, ..., p_L^n)$, $V^n(S^n)$ as the current largest estimate, i.e, $\max_{x\in\mathcal{X}} \bar{f}^n(x)$ (recall that $\bar{f}^n(x)$ is our estimate of $f(x;\theta^*)$ at time n, given by $\bar{f}^n(x) = \sum_{i=1}^L f(x;\theta_i^n)p_i^n$). Let $p^{n+1}(x)$ represent the posterior probability after measuring x. KGDP-f is calculated in [7] as:

$$\begin{aligned}\nu^{KGDP-f,n}(x) =& \mathbb{E}^n\left[\max_{x'}\sum_{i=1}^L f_i(x')p_i^{n+1}(x)|S^n = s, x^n = x\right] - \max_{x'}\sum_{i=1}^L f_i(x')p_i^n \\ =& \sum_{j=1}^L\left[\int_\omega \max_{x'}\frac{1}{c_j}\sum_{i=1}^L f_i(x')\exp\left(-\frac{[f_j(x)-f_i(x)+\omega]^2}{2\sigma^2}\right)p_i^n g(\omega)d\omega\right]p_j^n \\ & - \max_{x'}\sum_{i=1}^L f_i(x')p_i^n, \end{aligned} \tag{3.5}$$

where i and j index the candidates in $\mathbb{L}^n$, $c_j = \sum_{i=1}^L \exp\left[-\frac{(f_j(x)-f_i(x)+\omega)^2}{2\sigma^2}\right]p_i^n$, $g(\omega) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{\omega^2}{2\sigma^2}\right)$, and $f_i(x)$ is short for $f(x;\theta_i^n)$.

In each integral, let $f_j(x) + \omega = \hat{y}$, and the above equation can be simplified as

$$\begin{aligned}\nu^{KGDP-f,n}(x) =& \frac{1}{\sqrt{2\pi}\sigma}\int_{-\infty}^{+\infty}\max_{x'}\left[\sum_{i=1}^L f_i(x')p_i^n\exp\left(-\frac{[\hat{y}-f_i(x)]^2}{2\sigma^2}\right)\right]d\hat{y} \\ & - \max_{x'}\sum_{i=1}^L f_i(x')p_i^n, \end{aligned} \tag{3.6}$$

where i indexes the candidates at time n. Compared with Equation (3.5) in [7], Equation (3.6) is simpler for both calculation and theoretical proof.

In KGDP-f, our decision at time n is

$$x^n = \underset{x\in\mathcal{X}}{\operatorname{argmax}}\,\nu^{KGDP-f,n}(x).$$

3.5.2. KGDP-H. Entropy is a metric to measure the uncertainty of unknown data, which is widely used especially in information theory. The entropy of the candidates at time n is given by

$$H(p_1^n, ..., p_L^n) = -\sum_{i=1}^L p_i^n \log p_i^n.$$

In KGDP-H, we measure the alternative that has the largest expected entropic loss. Define the state variable S^n as the probability vector $(p_1^n, p_2^n, ..., p_L^n,)$, and let $V^n(S^n)$ be the entropy. The KGDP-H score is given by

$$\begin{aligned}\nu^{KGDP-H,n}(x) &= \mathbb{E}^n\left[\sum_{i=1}^L p_i^{n+1}(x)\log p_i^{n+1}(x)|S^n = s, x^n = x\right] - \sum_{i=1}^L p_i^n\log p_i^n \\ &= \sum_{j=1}^L\left(\int_\omega\sum_{i=1}^L p_{i|j}^{n+1}(x,\omega)\log p_{i|j}^{n+1}(x,\omega)g(\omega)d\omega\right)p_j^n - \sum_{i=1}^L p_i^n\log p_i^n, \end{aligned} \tag{3.7}$$

where $g(\omega) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\omega^2}{2\sigma^2}\right)$, and $p_{i|j}^{n+1}(x,\omega)$ is the probability of θ_i at time $(n+1)$ given that θ_j is the truth and the noise is ω, given by

$$p_{i|j}^{n+1}(x,\omega) = \frac{\exp\left[-\frac{(f_j(x)-f_i(x)+\omega)^2}{2\sigma^2}\right] p_i^n}{\sum_{k=1}^{L} \exp\left[-\frac{(f_j(x)-f_k(x)+\omega)^2}{2\sigma^2}\right] p_k^n}.$$

Alternatively, it can be written as

$$\begin{aligned}\nu^{KGDP-H,n}(x) =& \int_{-\infty}^{+\infty} \sum_{i=1}^{L} p_i^{n+1}(x) \log p_i^{n+1}(x) \cdot \frac{1}{\sqrt{2\pi}\sigma} \left(\sum_{i=1}^{L} p_i^n \exp\left[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}\right] \right) d\hat{y} \\ & - \sum_{i=1}^{L} p_i^n \log p_i^n \\ =& \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} \sum_{i=1}^{L} p_i^n \exp\left[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}\right] \log p_i^{n+1}(x) d\hat{y} - \sum_{i=1}^{L} p_i^n \log p_i^n, \end{aligned} \tag{3.8}$$

where $p_i^{n+1}(x)$ is given by Equation (3.2).

In KGDP-H, our decision at time n is

$$x^n = \underset{x\in\mathcal{X}}{\operatorname{argmax}}\, \nu^{KGDP-H,n}(x).$$

After we exhaust the budget of N experiment, we give our estimates of the optimal alternative $\hat{x}^*$ and parameters $\hat{\theta}^*$ by:

$$\hat{x}^* = \underset{x\in\mathcal{X}}{\operatorname{argmax}}\, \bar{f}^N(x) = \underset{x\in\mathcal{X}}{\operatorname{argmax}} \sum_{i=1}^{L} p_i^N f(x;\theta_i^N), \qquad \hat{\theta}^* = \underset{\theta\in\mathbb{K}}{\operatorname{argmin}}\, MSE(\theta|\mathcal{F}^N).$$

3.6. Detailed Flowchart. The detailed flow of the whole procedure is shown in Algorithm 1, where the function *select_x*(*policy*, ...) uses the designated *policy* to select an alternative to measure. The flowchart of our experiment decision *select_x*() is shown in Algorithm 2. In *select_x*(), besides the KGDP-f and KGDP-H policies, we also include three competing policies used in our experiments, namely (1) Pure Exploration, which chooses x^n randomly; (2) Pure Exploitation, which always chooses the current best alternative; and (3) Max Variance (Max-Var), which picks the alternative that has the largest variance under prior belief.

4. Convergence of KGDP without Resampling. By construction, KGDP-f and KGDP-H are naturally the optimal myopic policy to find the optimal alternative and parameter (in terms of entropy) respectively. In this section, we show that KGDP-f and KGDP-H with truth from prior are asymptotically optimal in the non-resampling scenario, a result that was not shown in [7] where KGDP was first introduced. In other words, as the budget $N \to \infty$, KGDP-f and KGDP-H will both find the optimal alternative and the correct parameters.

We denote the finite set of alternatives as $\mathcal{X}$. Define $(\Omega, \mathcal{F}, \mathbb{P})$ as the probability space, where Ω is the set of all possible measurement histories $\{(x^0, \hat{y}^1), ..., (x^{N-1}, \hat{y}^N)\}$, and $\mathcal{F}$ is the σ-algebra generated by the history. In this section, $N = \infty$.

To prove asymptotic optimality, we first show that if an alternative x is measured infinitely often, we can learn its true function value almost surely. As N goes to infinity, there will be a subset of alternatives that are evaluated infinitely often. We prove that under KGDP-f or KGDP-H policy with truth from prior, $f(x, \theta^*)$ is the

Algorithm 1 Flow of the resampling algorithm

Input: Budgets: N; Alternatives: $\mathcal{X} = \{x_1, ..., x_M\}$; Noise: σ; Large pool: $\mathbb{K}$; Smaller pool size: R; Number of candidates: L;; Resample stepsize: n^{resamp}; Threshold for probability: ϵ; and the *policy* to select x^n.

Output: Estimate of the optimal alternative: $\hat{x}^*$, estimate of parameters: $\hat{\theta}^*$.

Choose L samples out of $\mathbb{K}$ randomly.
Set $p_1^0 = ... = p_L^0 = \frac{1}{L}$.
for $n = 0$ to $N - 1$ **do**
  $x^n = select_x(policy, \mathcal{X}, \vec{p^n}, \mathbb{L}^n)$.
  Take a measurement: $\hat{y}^{n+1}$.
  Update each p_i using:

$$p_i^{n+1} = \frac{\exp[-\frac{(\hat{y}^{n+1} - f_i(x))^2}{2\sigma^2}]}{\sum_{l=1}^{L} \exp[-\frac{(\hat{y}^{n+1} - f_l(x))^2}{2\sigma^2}] p_l^n} p_i^n.$$

  if $n + 1 \equiv 0 \bmod (n^{resamp})$ **or** $|\{p_i^{n+1} \geq \epsilon\}| \geq L/2$ **then**
    Calculate MSE for all the K samples.
    Construct $\mathbb{S}^n$ as the set of R θ's with the smallest MSE.
    Calculate the likelihood of each θ in $\mathbb{S}^n$.
    Construct $\mathbb{L}_{rm}^n$ as the set of θ's to be removed, and set $p_i^{n+1} = 0$ for $\theta_i \in \mathbb{L}_{rm}^n$.
    while $\min(p^{n+1}) \leq \epsilon$ **do**
      Find candidates with $p^{n+1} \leq \epsilon$ and remove them.
      Select $|\mathbb{L}_{rm}^n|$ θ's from $\mathbb{S}^n$ by weighted sampling without replacement according to their likelihoods.
      Update posterior probabilities for all the L new θ's:

$$p_i^{n+1} = \frac{\prod_{j=0}^{n} \exp[-\frac{(\hat{y}^{j+1} - f(x^j; \theta_i^{n+1}))^2}{2\sigma^2}]}{\sum_{l=1}^{L} \prod_{j=0}^{n-1} \exp[-\frac{(\hat{y}^{j+1} - f(x^j; \theta_l^{n+1}))^2}{2\sigma^2}]}.$$

    end while
  end if
end for
return $\hat{x}^* = \text{argmax}_{x \in \mathcal{X}} \sum_{i=1}^{L} p_i^N f(x; \theta_i^N)$, $\hat{\theta}^* = \text{argmin}_{\theta \in \mathbb{K}} MSE(\theta | \mathcal{F}^N)$.

only function that fits the true function values of this subset. To show this, we reveal the nonnegativity of KGDP-f and KGDP-H scores, and the fact that any alternative measured infinitely often has its KGDP-f and KGDP-H scores converge to 0. Then we use proof by contradiction. Assume $f(x, \theta^*)$ is not the only function inferred by the infinitely measured alternatives. Then there exists at least an x measured for a finite number of times. We claim that it has positive KGDP-f and KGDP-H scores in the limit. This is contrary to the fact that either KGDP-f or KGDP-H policy chooses the alternative with the largest score.

We begin by showing that if we measure a point x infinitely often, the correct function value at x will be found almost surely:

LEMMA 4.1. *Let $N^n(x)$ be the number of measurements taken on x when the total number of measurements is n. If $N^n(x) \to \infty$ as $n \to \infty$, then for $\forall l \in \{1, 2, ..., L\}$ such that $f(x; \theta_l) \neq f(x; \theta^*)$, $p_l^n \to 0$ almost surely, i.e, $\mathbb{P}(\lim_{n \to \infty} p_l^n = 0) = 1$.*

The rigorous proof of Lemma 4.1 comes from the strong law of large numbers (shown in Appendix).

Algorithm 2 Choose the next alternative x^n.

Input: Policy: *Pure Exploration, Pure Exploitation, Max-Var, KGDP-f, KGDP-H*; Alternatives: $\mathcal{X} = \{x_1, ..., x_M\}$; Probability vector $\vec{p^n} = \{p_1^n, ..., p_L^n\}$; L candidates $\mathbb{L}^n = \{\theta_1^n, ..., \theta_L^n\}$.
Output: The alternative to measure next: x^n.

switch (*policy*)
case KGDP-f**:**
  $x^n = \text{argmax}_{x\in\mathcal{X}}\, \nu^{KGDP-f,n}(x)$.
case KGDP-H**:**
  $x^n = \text{argmax}_{x\in\mathcal{X}}\, \nu^{KGDP-H,n}(x)$.
case Pure Exploration**:**
  $x^n = rand(\{x_1, ..., x_M\})$.
case Pure Exploitation**:**
  $x^n = \text{argmax}_{x\in\mathcal{X}}\, \bar{f}^n(x)$.
case Max-Var**:**
  $x^n = \text{argmax}_{x\in\mathcal{X}} \sum_{j=1}^{L} p_j[f(x;\theta_j^n) - \bar{f}^n(x)]^2$.
end switch
return x^n.

We define the set of the almost sure events as Ω_0. For any $\omega \in \Omega$, we denote the set of alternatives as $\mathcal{X}_\infty(\omega)$ that are measured infinitely often. Hence, for any $\omega \in \Omega_0$ and any $x \in \mathcal{X}_\infty(\omega)$, we have $p^n(\theta^*)(\omega) \to 1$.

We hope that we can learn θ^* via $\mathcal{X}_\infty$. To achieve this goal, we first study some properties of KGDP-f and KGDP-H scores. We start by showing that the value of information for both objectives is always nonnegative:

LEMMA 4.2. *For $\forall n \geq 0, \forall x \in \mathcal{X}$, the KGDP-$f$ score $\nu^{KGDP-f,n}(x) \geq 0$. Equality holds if and only if (1) there exists x' such that $x' \in argmax_x\, f(x;\theta_i)$ for all i such that $p_i^n > 0$, or (2) all functions with $p^n > 0$ have the same value at x.*

LEMMA 4.3. *For $\forall n \geq 0, \forall x \in \mathcal{X}$, the KGDP-$H$ score $\nu^{KGDP-H,n}(x) \geq 0$. Equality holds if and only if all functions with $p^n > 0$ have the same value at x.*

Proof. [Sketch of proof for Lemma 4.2 and 4.3] (See Appendix for full proof.) First, we show $\mathbb{E}^n\left[p_i^{n+1}(x)\right] = p_i^n$. Then, applying Jensen's inequality will give us the nonnegativity of both KGDP-f and KGDP-H. □

According to Lemma 4.2 and Lemma 4.3, KGDP-f equals 0 in only two cases: (1) either when the functions are aligned, in which case $\nu^{KGDP,n}(x) = 0$ for all x, (2) or when the functions with nonzero probabilities have the same value at x, in which case $\nu^{KGDP,n}(x) = 0$ for this particular x. However, KGDP-H equals 0 only in the second case.

We then show that for any x measured infinitely often, its KGDP-f and KGDP-H scores also go to zero:

LEMMA 4.4. *For $\forall \omega \in \Omega_0$ and $x \in \mathcal{X}_\infty(\omega)$, we have $\lim_{n\to\infty} \nu^{KGDP-f,n}(x)(\omega) = 0$, and $\lim_{n\to\infty} \nu^{KGDP-H,n}(x)(\omega) = 0$.*

Intuitively, as we measure x infinitely often and learn the true value gradually, the condition for both KGDP-f and KGDP-H being 0 is satisfied (by Lemma 4.2 and 4.3). The full proof, as shown in Appendix, reveals the fact that given a fixed x, $\nu^{KGDP-f}(x)$ and $\nu^{KGDP-H}(x)$ are both uniformly continuous in the probability vector space $\vec{p^n} = (p_1^n, ..., p_L^n)$. Combined with the convergence result of p_i^n shown by Lemma 4.1, we conclude $\nu^{KGDP-f}(x)$ and $\nu^{KGDP-H}(x)$ also converge to 0.

For any subset of alternatives, if it is sufficient to infer θ^* by fitting the true function values of the subset, we call it a *sufficient set*:

DEFINITION 4.5. *We define a sufficient set $\mathcal{X}_s \subset \mathcal{X}$ as follows: for $\forall \theta \neq \theta^*$, there exists $x \in \mathcal{X}_s$ such that $f(x;\theta) \neq f(x;\theta^*)$.*

In other words, a sufficient set $\mathcal{X}_s$ is a subset of $\mathcal{X}$ upon which we can distinguish θ^* from all the others. Alternatively, for any subset $\mathcal{X}_s = \{x_1, ..., x_m\} \subset \mathcal{X}$, if $f(x;\theta^*)$ is the only function that fits the m points $(x_1, f(x_1;\theta^*)), ..., (x_m, f(x_m;\theta^*))$, then $\mathcal{X}_s$ is a sufficient set. Obviously, the largest sufficient set is $\mathcal{X}$. An example of an *insufficient set* is that, if all functions have the same values at $\{x_1, ..., x_m\}$, then this is not a sufficient set.

We show that under KGDP-f or KGDP-H policy, we measure a sufficient set almost surely:

LEMMA 4.6. *For any $\omega \in \Omega_0$, the alternatives measured infinitely often under the KGDP-f or KGDP-H policy constitute a sufficient set.*

Proof. [Sketch of proof] (See full proof in Appendix.) If $\mathcal{X}_\infty(\omega)$ is not a sufficient set, then at least another function other than $f(x;\theta^*)$ fits $\mathcal{X}_\infty(\omega)$. Hence, there exists $x \notin \mathcal{X}_\infty(\omega)$ such that this function has different values from $f(x;\theta^*)$. We can prove that for this x, $\lim_{n\to\infty} \nu^{KGDP-f,n}(x)(\omega) > 0$, and $\lim_{n\to\infty} \nu^{KGDP-H,n}(x)(\omega) > 0$. This is contrary to the fact that KGDP-f or KGDP-H will not measure any x out of $\mathcal{X}_\infty(\omega)$ after a finite number of times. □

Using the previous lemmas, we can conclude with the following asymptotic optimality result:

THEOREM 4.7. *Non-resampling KGDP-f with truth from prior is asymptotically optimal in finding both the optimal alternative and the correct parameter. The same holds for KGDP-H.*

The detailed proofs of Lemma 4.2 - 4.6 and Theorem 4.7 are given in Appendix.

5. Convergence of KGDP with Resampling. We now establish the important result that the resampling algorithm converges to the true optimal solution, both in terms of identifying the right parameter θ^* as well as the optimal design x^*.

As we move from the non-resampling case to the resampling one, the difference is that the candidates are no longer fixed, and θ^* is not guaranteed to be one of the candidates.

Note that the nonnegativity of the knowledge gradient established in Lemma 4.2 and 4.3 still holds in the resampling case, regardless of whether the candidates include the truth or not.

Unlike the non-resampling case, we have two probability spaces here. One is still $(\Omega, \mathcal{F}, \mathbb{P})$ as above, defined on the measurement history, while the other one describes the randomness of selecting candidates in the resampling process. We denote the latter as $(\Omega_\theta, \mathcal{F}_\theta, \mathbb{P}_\theta)$, where Ω_θ is the set of all possible selection combinations of the L candidates in N iterations (in this section, $N = \infty$). The full probability space, denoted as $(\Omega_f, \mathcal{F}_f, \mathbb{P}_f)$, is the product space of $(\Omega, \mathcal{F}, \mathbb{P})$ and $(\Omega_\theta, \mathcal{F}_\theta, \mathbb{P}_\theta)$. We use ω to represent an element in Ω and ω_θ in Ω_θ.

Let $\mathbb{K}$ denote the large pool, i.e, the set of the K θ's. We can adapt the concept of a sufficient set by considering the entire set $\mathbb{K}$. For a subset of alternatives $\mathcal{X}_s \subset \mathcal{X}$, if $f(x;\theta^*)$ is the only function among all the K functions that fits the true values of alternatives in $\mathcal{X}_s$, we call it a sufficient set.

Unlike the non-resampling case, since the candidates may keep changing, the limit of a particular p_i^n may not exist. However, if we regard the whole $\mathbb{K}$ as the candidate set and consider the probability of each θ, then Lemma 4.1 still holds, demonstrating

that infinite measurements on x give its true value. We define Ω_1 as the set of almost sure events in Ω for Lemma 4.1 to hold.

LEMMA 5.1. *For any $\omega \in \Omega_1$, the alternatives measured infinitely often under the KGDP-f or KGDP-H policy constitute a sufficient set. We denote this set as $\mathcal{X}_\infty(\omega)$.*

Proof. [Sketch of proof] (See Appendix for full proof.) Assume the contrary. Then there exists at least one $\theta \neq \theta^*$, such that $f(x;\theta) = f(x;\theta^*)$ for any $x \in \mathcal{X}_\infty(\omega)$. Denote the set of such θ's as $\mathbb{K}'$. We can show that for this fixed ω, the event happens infinitely often that at least two θ's out of $\mathbb{K}' \bigcup \{\theta^*\}$ are included in the candidate set. At these times, we can further show that there exists an $x \notin \mathcal{X}_\infty(\omega)$ and $\epsilon > 0$ such that $\nu^{KGDP-f,n}(x) > \epsilon$ and $\nu^{KGDP-H,n} > \epsilon$, while on the other hand, for any $x \in \mathcal{X}_\infty(\omega)$, $\lim_{n\to\infty} \nu^{KGDP-f,n}(x) = 0$ and $\lim_{n\to\infty} \nu^{KGDP-H,n}(x) = 0$. This is contradictory to the fact that any $x \notin \mathcal{X}_\infty(\omega)$ is measured for a finite number of times. □

Lemma 5.1 implies the following theorem, which is also our main result in this section:

THEOREM 5.2. *Resampling KGDP-f is asymptotically optimal in finding the optimal alternative and the correct parameter. The same holds for KGDP-H.*

Proof. [Sketch of proof] (See Appendix for full proof.) The crucial step is to show that for any $\omega \in \Omega_1$, there exists time $T(\omega)$, such that θ^* always appears in the candidate set after $T(\omega)$ with probability 1. Note that this probability is calculated in the $(\Omega_\theta, \mathcal{F}_\theta, \mathbb{P}_\theta)$ space. Extending this result to the full probability space $(\Omega_f, \mathcal{F}_f, \mathbb{P}_f) = (\Omega, \mathcal{F}, \mathbb{P}) \times (\Omega_\theta, \mathcal{F}_\theta, \mathbb{P}_\theta)$, we can see that $\lim_{n\to\infty} p^n(\theta^*) = 1$ happens with probability 1. □

6. Empirical Results. We study an asymmetric unimodal problem as a benchmark function, and another material science problem, known as the nanoemulsion problem in a water-oil-water (W/O/W) model.

We mainly demonstrate results from the following aspects:

- The visualization of the resampling process, including the evolution of the candidates, small pool, and approximated function $\bar{f}^n(x)$;
- The performance of resampling KGDP-f and KGDP-H, both in relation to each other and their absolute performance in terms of different metrics;
- The relative performance of resampling KGDP-f and KGDP-H policies versus competing policies;
- The comparison of resampling and non-resampling KGDP methods;
- The empirical rate of convergence under different noise and dimensionality settings.

We feel that the rate of convergence result is the most important.

6.1. An Asymmetric Unimodal Benchmark Function. We study a multi-dimensional benchmark function, which produces a wide range of asymmetric, unimodular functions with heteroscedastic noise, which we have found to be a challenging class of test problems. The function is given by

$$f(x_1, ..., x_k) = \sum_{i=1}^{k} \eta_{1,i} \mathbb{E}[\min(x_i, (D - \sum_{j=1}^{i-1} x_j)^+)] - \sum_{i=1}^{k} \eta_{2,i} x_i,$$

where $x_1, ..., x_k$ are the decision vector, D is a uniformly distributed random variable, $\eta_{1,i}$ are fixed constants, and $\eta_{2,i}$ are unknown parameters. In the following experiments, we fix the variance of D, but set its mean as well as $\eta_{2,i}$'s as parameters (θ), which results in a $(k+1)$-dimensional problem.

6.1.1. Illustration of the Resampling Process. We first show the results of resampling KGDP-f in two-dimensional cases, where we can plot the θ space.

Figure 6.1 shows the evolution of the small pool and the L candidates in θ space in the first several iterations of a realization of the problem. In these images, the horizontal and vertical axes represent θ_1 and θ_2 respectively, with color indicating the values of MSE. The red pentagon shows where the truth θ^* is located, while the white star indicates the θ with the smallest MSE in the large pool. The region circled by the red line is the range of the small pool. The green squares indicate the locations of the candidates, and the sizes are proportional to their probabilities. Note that these images have been scaled to show more detail. We can see that the small pool shrinks quickly towards around θ^*, and the θ with the smallest MSE as well as the candidates also converges to θ^* within only a few iterations.

6.1.2. Comparison of Different Policies to Choose x^n. We use the opportunity cost (OC) to evaluate the performance of various policies from the alternative perspective. OC is defined as the difference between the values at our estimated optimal x and the true optimal x, i.e.,

$$OC(n) = \max_{x\in\mathcal{X}} f(x;\theta^*) - f(\operatorname*{argmax}_{x\in\mathcal{X}} \bar{f}^n(x);\theta^*).$$

To normalize the opportunity cost, we define as the percentage OC the ratio with respect to the optimal function value, i.e.:

$$OC\%(n) = \frac{\max_{x\in\mathcal{X}} f(x;\theta^*) - f(\operatorname{argmax}_{x\in\mathcal{X}} \bar{f}^n(x);\theta^*)}{\max_{x\in\mathcal{X}} f(x;\theta^*)}.$$

We define the noise level as the ratio of the noise standard deviation to the range of the true function, i.e, $\frac{\sigma}{\max_x f(x;\theta^*) - \min_x f(x;\theta^*)}$. When we say the noise is 20%, we mean this ratio is 20%.

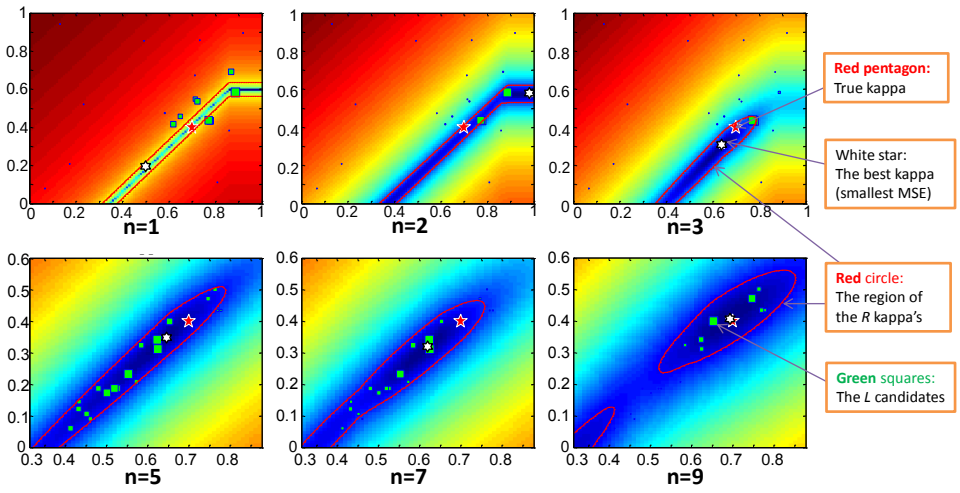

Fig. 6.1: Illustration of resampling in θ space (Iterations 1, 2, 3, 5, 7, 9) (5,7,9 are zoomed in).

Figure 6.2 shows the comparison of the five policies in three dimensional settings. The subimages correspond to 5%, 20% and 50% of noise from left to right. We can see that KGDP-f and KGDP-H are among the best under different noise levels. We have similar results for two- and five-dimensional cases, too.

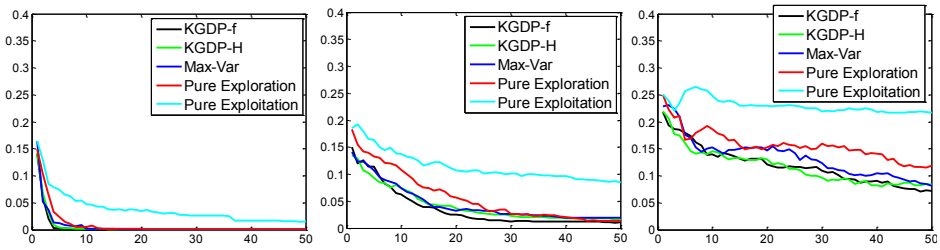

Fig. 6.2: OC% of different policies for three-dimensional θ (5%, 20% and 50% noise from left to right).

6.2. Application in W/O/W Nanoemulsion Stability Study. We study the same problem of water-oil-water(W/O/W) nanoemulsion stability as in [7], and repeat its experiments using the original non-resampling KGDP and our resampling KGDP methods to compare the results.

The problem studies controlled payload delivery using water-oil-water (W/O/W) double nanoemulsions, which depicts the delivery of payload molecules from the internal water droplet, surrounded by a larger droplet of oil, to the external water. This process is performed via laser excitation and controlled by several experimental variables. We can model the stability of this process as a function $f(x;\theta)$ that is nonlinear in θ (as well as x). x is a five-dimensional control variable and includes variables such as the initial volume of the external water, the initial volume fraction of the oil droplets, and the diameter of the oil droplets. θ includes seven dimensions, representing the unknown parameters that appear in the formulation, such as the energy barrier for flocculation, the rate prefactor for coalescence and droplet adsorption/desorption energy barrier difference. Our goal is to conduct a series of experiments to learn the optimal setting (x) to achieve the best stability, and also to learn the correct unknown parameters θ. A more detailed introduction and formulation of this problem can be found in [7].

In order to be consistent with [7], we fix three dimensions of x and create a finite set of alternatives using the other two dimensions. For θ, we study 1) a three dimensional case and 2) a seven dimensional one in our experiments. Remember that the dimensionality of θ is more important than the dimensionality of x, since we always have finite samples of x.

6.2.1. Three Dimensional θ. Figure 6.3 shows the comparison of the various policies in terms of OC% under 50% of noise. The left image is a comparison across the five policies, while the right one compares the resampling method with non-resampling KGDP in [7].

The left image indicates that Max-Var, KGDP-f and KGDP-H are among the best, showing very close performance. The right image compares: (1) the orange curve: the non-resampling KGDP-f with truth from prior, which is the ideal case; (2) the pink curve: the non-resampling KGDP-f with truth not from prior, which is

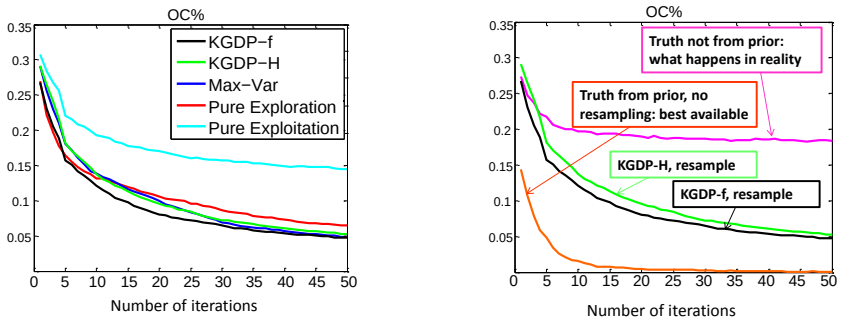

(a) Comparison across five policies

(b) Comparison with non-resampling KGDP

Fig. 6.3: OC% of different policies in three-dimensional nanoemulsion (50% noise).

what happens with truth not from prior using [7]; (3) the black curve: resampling KGDP-f; and (4) the green curve: resampling KGDP-H. We can see that (2) gets stuck after a few iterations, while our resampling methods using either KGDP-f or KGDP-H are both able to find the more probable parameters and approach the ideal case gradually.

We have shown the performance of our resampling algorithm in finding the best alternative in terms of OC%. To show its performance in learning the parameters, we denote our best estimate of θ^* as $\hat{\theta}^*$, and calculate:

(1) the mean square error of $f(x;\theta^*)$ and $f(x;\hat{\theta}^*)$, defined as (remember M is the size of $\mathcal{X}$):

$$f_{MSE} = \frac{1}{|\mathcal{X}|}\sum_{x\in\mathcal{X}} |f(x;\theta^*) - f(x;\hat{\theta}^*)|^2,$$

(2) or the error of a particular dimension of θ, say $|\theta_1^* - \hat{\theta}_1^*|$.

Note that in method (2), we do not calculate the error of all dimensions as $||\theta^* - \hat{\theta}^*||$. This is because different dimensions of θ may have different units and follow various distributions, hence it is meaningless to sum up the errors of different dimensions. Moreover, in a higher dimensional problem, some dimensions may converge more quickly than others, and while the entire vector converges in aggregate, individual dimensions may converge more slowly than others.

We show the MSE and error of θ in Figure 6.4. In the first image, we calculate f_{MSE}, and plot $\frac{\sqrt{f_{MSE}}}{\max_x f(x;\theta^*) - \min_x f(x;\theta^*)}$. and the other three show $|\theta_1^* - \hat{\theta}_1^*|$, $|\theta_2^* - \hat{\theta}_2^*|$ and $|\theta_3^* - \hat{\theta}_3^*|$, respectively. Note that in this example, the second dimension seems less important than the other two. We also have similar results under 20% of noise.

6.2.2. Seven Dimensional θ. In a seven-dimensional case, Figure 6.5 shows the process of $\bar{f}^n(x)$ approaching the true function, using resampling KGDP-f under 20% noise. The four images in the second row are our approximations at $n = 1, 10, 20, 45$ respectively.

The result of OC% under 50% noise are shown in Figure 6.6. We have similar results regarding learning the parameters as in three-dimensional cases.

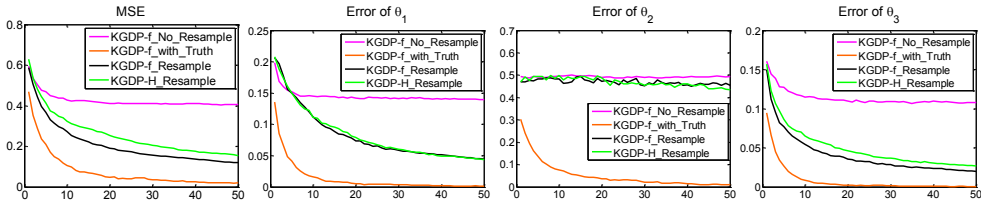

Fig. 6.4: Performance of resampling KGDP in finding θ in three-dimensional nanoemulsion (50% noise).

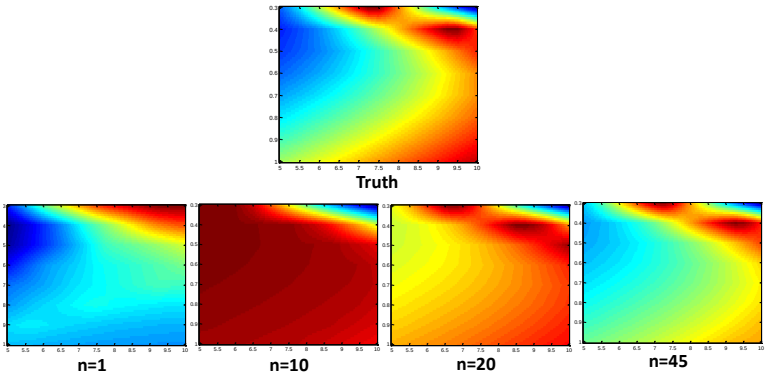

Fig. 6.5: Evolution of $\bar{f}^n(x)$ (seven-dimensional θ, 20% noise).

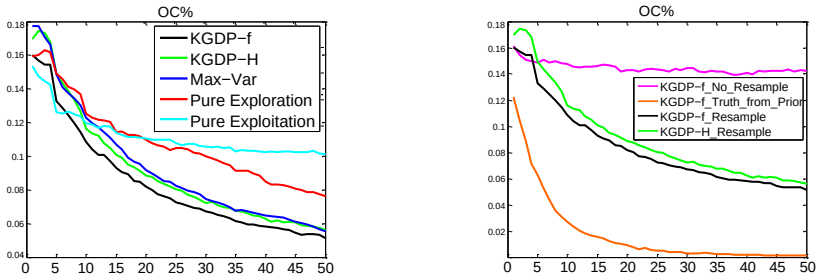

(a) Comparison across five policies (b) Comparison with non-resampling KGDP

Fig. 6.6: OC% of different policies in seven-dimensional nanoemulsion (50% noise).

7. Conclusion. Optimal learning for nonlinear belief models has been an interesting but difficult problem. With a belief model nonlinear in unknown parameters, the knowledge gradient becomes computationally intractable if using continuous parameter space. In [7], the Knowledge Gradient with Discrete Priors (KGDP) model is proposed, which uses a sampled representation of the parameter space. This method works well when one of the sampled candidates is exactly the truth, but this would

not happen in practice.

In this paper, we propose a resampling algorithm that cooperates with KGDP to solve this problem. We no longer require that one of the candidates is correct, but use the measurement history to guide resampling and discovering more probable parameters in the parameter space. Within a small budget of noisy and expensive experiments, our algorithm is capable of both finding the optimal alternative that maximizes the value function, and narrow down the location of the unknown parameters. Experiments on both a synthetic benchmark function and a real world materials science problem, known as the stability of W/O/W nanoemulsion problem, have shown strong performance of our algorithm.

We recognize the duality of this optimization problem and manipulate in both the alternative and the parameter spaces. Besides the idea of maximizing the value function (denoted as KGDP-f) as used in previous optimal learning literatures, we also put forward an entropy-oriented metric that works under the knowledge gradient framework, which aims directly at lowering the uncertainty of the parameters (in terms of entropy). Experiments show strong performance of both metrics.

We also prove the asymptotic optimality of KGDP-f and KGDP-H in both the non-resampling with truth from prior case, and the resampling framework. That means, given an infinite number of measurements, we are able to find both the optimal alternative and the correct parameter. Hence, KGDP-f and KGDP-H are both myopically and asymptotically optimal in learning the alternative and reducing parameter uncertainty respectively.

Appendix A.

Lemma 4.1. Let $N^n(x)$ be the number of measurements taken on x when the total number of measurements is n. If $N^n(x) \to \infty$ as $n \to \infty$, then for $\forall l \in \{1, 2, ..., L\}$ such that $f(x; \theta_l) \neq f(x; \theta^*)$, $p_l^n \to 0$ almost surely, i.e, $\mathbb{P}(\lim_{n\to\infty} p_l^n = 0) = 1$.

Proof of Lemma 4.1. We prove the following claim first.

Claim 1: $\forall m \in \{1, ..., M\}$, let $p_1^n, p_2^n, ..., p_L^n$ be the probabilities of the L θ's after n measurements at x_m (i.e, we take n measurements only at x and nowhere else). For $\forall \epsilon > 0$, $\forall l \in \{1, 2, ..., L\}$ such that $f(x_m; \theta_l) \neq f(x_m; \theta^*)$, $\mathbb{P}(p_l^n < \epsilon) = 1$, as $n \to \infty$.

Proof of Claim 1:

If we take n measurements at x_m, suppose the results are $\{\hat{y}^1, ..., \hat{y}^n\}$. Let $f_l = f(x_m; \theta_l)$. Then $p_1, ..., p_L$ are given by:

$$p_l^n = \frac{\prod_{j=1}^n \exp[-\frac{(\hat{y}^j - f_l)^2}{2\sigma^2}]}{\sum_{i=1}^L \prod_{j=1}^n \exp[-\frac{(\hat{y}^j - f_i)^2}{2\sigma^2}]}.$$

Let $f^* = f(x_m; \theta^*)$. Then $\hat{y}^j \sim N(f^*, \sigma^2)$. Let $Z^j = \frac{\hat{y}^j - f^*}{\sigma}$, then $\hat{y}^j \sim f^* + \sigma Z^j$, $Z^j \sim N(0, 1)$.

For any l such that $f_l \neq f^*$,

$$\begin{aligned} p_l^n \leq & \frac{\prod_{j=1}^n \exp[-\frac{(\hat{y}^j - f_l)^2}{2\sigma^2}]}{\prod_{j=1}^n \exp[-\frac{(\hat{y}^j - f^*)^2}{2\sigma^2}]} = \frac{\exp[-\frac{\sum_{j=1}^n (\hat{y}^j - f_l)^2}{2\sigma^2}]}{\exp[-\frac{\sum_{j=1}^n (\hat{y}^j - f^*)^2}{2\sigma^2}]} = \frac{\exp[-\frac{\sum_{j=1}^n (f^* - f_l + \sigma Z^j)^2}{2\sigma^2}]}{\exp[-\frac{\sum_{j=1}^n (Z^j)^2}{2}]} \\ = & \exp\left(-\frac{\sum_{j=1}^n [(f^* - f_l + \sigma Z^j)^2 - (\sigma Z^j)^2]}{2\sigma^2}\right) \\ = & \exp\left(-\frac{n(f^* - f_l)^2 + 2\sigma(f^* - f_l)\sum_{j=1}^n Z^j}{2\sigma^2}\right). \end{aligned} \tag{A.1}$$

For any $\epsilon > 0$, $\exp\left(-\frac{n(f^*-f_l)^2+2\sigma(f^*-f_l)\sum_{j=1}^n Z^j}{2\sigma^2}\right) < \epsilon \Rightarrow p_l^n < \epsilon$, and therefore $\{\lim_{n\to\infty} \exp\left(-\frac{n(f^*-f_l)^2+2\sigma(f^*-f_l)\sum_{j=1}^n Z^j}{2\sigma^2}\right) = 0\} \subseteq \{\lim_{n\to\infty} p_l^n = 0\}$. We show that

$$\mathbb{P}\left[\lim_{n\to\infty} \exp\left(-\frac{n(f^*-f_l)^2+2\sigma(f^*-f_l)\sum_{j=1}^n Z^j}{2\sigma^2}\right) = 0\right] = 1,$$

which implies the almost sure convergence of p_l^n.

By the strong law of large numbers, $\frac{\sum_{j=1}^n Z^j}{n} \xrightarrow{a.s.} 0$. That is, there exists $\Omega_{x,l} \in \Omega$, such that for $\forall \omega \in \Omega_{x,l}$, $\lim_{n\to\infty} \frac{\sum_{j=1}^n Z^j(\omega)}{n} = 0$, and $\mathbb{P}(\Omega_{x,l}) = 1$. Our goal is to show that for $\forall \epsilon > 0$, $\forall \omega \in \Omega_{x,l}$,

$$\lim_{n\to\infty} \exp\left(-\frac{n(f^*-f_l)^2+2\sigma(f^*-f_l)\sum_{j=1}^n Z^j(\omega)}{2\sigma^2}\right) = 0.$$

Without loss of generality, we assume $f^* - f_l \geq 0$. For $\forall \omega \in \Omega_{x,l}$, there exists $N_1 \in \mathbb{N}$ such that for $\forall n > N_1$, $\left|\frac{\sum_{j=1}^n Z^j(\omega)}{n}\right| < \frac{f^*-f_l}{4\sigma}$.

Therefore, for any $n > N_1$, $n\left[(f^*-f_l)^2 + 2\sigma(f^*-f_l)\frac{\sum_{j=1}^n Z^j(\omega)}{n}\right] > n \cdot \frac{(f^*-f_l)^2}{2}$. Hence, for $\forall n > N_1$,

$$\exp\left(-\frac{n(f^*-f_l)^2+2\sigma(f^*-f_l)\sum_{j=1}^n Z^j(\omega)}{2\sigma^2}\right) < \exp\left(-\frac{n(f^*-f_l)^2}{4\sigma^2}\right). \tag{A.2}$$

For $\forall \epsilon > 0$, there exists N_2 such that for $\forall n > N_2$, $\exp\left(-\frac{n(f^*-f_l)^2}{4\sigma^2}\right) < \epsilon$.

Take $N = \max(N_1, N_2)$, we know that for $\forall \omega \in \Omega_{x,l}$, $\forall \epsilon > 0$, there exists $N > 0$, such that for $\forall n > N$,

$$\exp\left(-\frac{n(f^*-f_l)^2+2\sigma(f^*-f_l)\sum_{j=1}^n Z^j(\omega)}{2\sigma^2}\right) < \epsilon.$$

Hence,

$$\Omega_{x,l} \subseteq \{\lim_{n\to\infty} \exp\left(-\frac{n(f^*-f_l)^2+2\sigma(f^*-f_l)\sum_{j=1}^n Z^j}{2\sigma^2}\right)\} \subseteq \{\lim_{n\to\infty} p_l^n = 0\}.$$

Since $\mathbb{P}(\Omega_{x,l}) = 1$, we have $\mathbb{P}\{\lim_{n\to\infty} p_l^n = 0\} = 1$. □

Back to Lemma 4.1, define $N^n(x,j)$ as the actual time when we make the j-th measurement on x ($1 \leq j \leq N^n(x)$). We have:

$$\begin{aligned}
p_l^n &= \frac{\prod_{j=1}^n \exp[-\frac{(\hat{y}^j - f(x^{j-1};\theta_l))^2}{2\sigma^2}]}{\sum_{i=1}^L \prod_{j=1}^n \exp[-\frac{(\hat{y}^j - f(x^{j-1};\theta_i)^2}{2\sigma^2}]} \leq \frac{\prod_{j=1}^n \exp[-\frac{(\hat{y}^j - f(x^{j-1};\theta_l))^2}{2\sigma^2}]}{\prod_{j=1}^n \exp[-\frac{(\hat{y}^j - f(x^{j-1};\theta^*)^2}{2\sigma^2}]} \\
&= \prod_{x\in\mathcal{X}} \frac{\prod_{j=1}^{N^n(x)} \exp[-\frac{(\hat{y}^{t^n(x,j)} - f(x^{j-1};\theta_l))^2}{2\sigma^2}]}{\prod_{j=1}^{N^n(x)} \exp[-\frac{(\hat{y}^{N^n(x,j)} - f(x^{j-1};\theta^*)^2}{2\sigma^2}]} \\
&\overset{def}{=\!=} \prod_{x\in\mathcal{X}} r(x) \\
&= \prod_{x\in\{x:f_l(x)=f^*(x)\}} r(x) \prod_{\substack{x\in\{x:N^n(x)=\infty,\\ f_l(x)\neq f(x;\theta^*)\}}} r(x) \prod_{\substack{x\in\{x:N^n(x)\neq\infty,\\ f_l(x)\neq f(x;\theta^*)\}}} r(x).
\end{aligned}$$

The first item equals 1, and the third equals a constant number. For any $\omega \in \bigcap_{x\in\mathcal{X}} \Omega_{x,l}$, the second item has limit as 0 by Claim 1. Since $\mathbb{P}(\bigcap_{x\in\mathcal{X}} \Omega_{x,l}) = 1 - \mathbb{P}(\bigcup_{x\in\mathcal{X}} \Omega^c_{x,l}) > 1 - \sum_{x\in\mathcal{X}} \mathbb{P}(\Omega^c_{x,l}) = 1$, we have $p^n_l \xrightarrow{a.s} 0$.

□

Lemma 4.2. For $\forall n \geq 0, \forall x \in \mathcal{X}$, the KGDP-$f$ score $\nu^{KGDP-f,n}(x) \geq 0$. Equality holds if and only if (1) there exists x' such that $x' \in \text{argmax}_x f(x;\theta_i)$ for all i such that $p^n_i > 0$, or (2) $p^n_i = 0$ if $f(x;\theta_i) \neq f(x;\theta^*)$.

Proof of Lemma 4.2. Suppose at time n, the probabilities of our prior candidates are $(p^n_1, p^n_2, ..., p^n_L)$. Since max is a convex function, by Jensen's inequality,

$$\begin{aligned}\nu^{KGDP,n}(x) &= \mathbb{E}^n[\max_{x'} \sum_{i=1}^{L} f_i(x')p^{n+1}_i(x)|S^n = s, x^n = x] - \max_{x'} \sum_{i=1}^{L} f_i(x')p^n_i \\ &\geq \max_{x'} \mathbb{E}[\sum_{i=1}^{L} f_i(x')p^{n+1}_i(x)|S^n = s, x^n = x] - \max_{x'} \sum_{i=1}^{L} f_i(x')p^n_i.\end{aligned}$$

We show that $\forall\ x,\ x',\ \mathbb{E}[\sum_{i=1}^{L} f_i(x')p^{n+1}_i(x)|S^n = s, x^n = x] = \sum_{i=1}^{L} f_i(x')p^n_i$. We know

$$\mathbb{E}[\sum_{i=1}^{L} f_i(x')p^{n+1}_i(x)|S^n = s, x^n = x] = \sum_{i=1}^{L} f_i(x')\mathbb{E}[p^{n+1}_i(x)|S^n = s, x^n = x].$$

Suppose the measurement at x is $\hat{y}$, we have

$$p^{n+1}_i(x) = \frac{p^n_i \exp[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}]}{\sum_{j=1}^{L} p^n_j \exp[-\frac{(\hat{y}-f_j(x))^2}{2\sigma^2}]}.$$

Based on our prior belief, the distribution of $\hat{y}$ is:

$$\hat{y} \sim \begin{cases} N(f_1(x;\theta_1), \sigma^2), \text{with probability } p^n_1 \\ N(f_2(x;\theta_2), \sigma^2), \text{with probability } p^n_2 \\ ... \\ N(f_L(x;\theta_L), \sigma^2), \text{with probability } p^n_L \end{cases}$$

Therefore,

$$\begin{aligned}\mathbb{E}\left[p^{n+1}_i(x)\right] &= \mathbb{E}\left(\frac{p^n_i \exp[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}]}{\sum_{j=1}^{L} p^n_j \exp[-\frac{(\hat{y}-f_j(x))^2}{2\sigma^2}]}\right) \\ &= \int_{-\infty}^{\infty} \sum_{k=1}^{L} p^n_k \frac{p^n_i \exp[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}]}{\sum_{j=1}^{L} p^n_j \exp[-\frac{(\hat{y}-f_j(x))^2}{2\sigma^2}]} \frac{1}{\sqrt{2\pi}} \exp[-\frac{(\hat{y}-f_k(x))^2}{2\sigma^2}] d\hat{y} \\ &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} p^n_i \exp[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}] \frac{\sum_{k=1}^{L} p^n_k \exp[-\frac{(\hat{y}-f_k(x))^2}{2\sigma^2}]}{\sum_{j=1}^{L} p^n_j \exp[-\frac{(\hat{y}-f_j(x))^2}{2\sigma^2}]} d\hat{y} \\ &= p^n_i \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} \exp[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}] d\hat{y} \\ &= p^n_i.\end{aligned}$$

Hence, for any x and x', $\mathbb{E}[\sum_{i=1}^{L} f_i(x')p_i^{n+1}(x)|S^n = s, x^n = x] = \sum_{i=1}^{L} f_i(x')p_i^n$.

For any convex function f, the general form of Jensen's inequality is $\mathbb{E}f(X) \geq f(\mathbb{E}X)$. The equality of Jensen's inequality requires that (1) f is linear on x, or (2) X is constant. (1) and (2) are equivalent to the two conditions stated in Lemma 4.2 respectively:

For (1), $\max(X)$ is not a linear function in general, except in the aligned case as stated in Lemma 4.2, Condition 1.

For (2), this is equivalent to the following statement: $\forall\ x'$, $\sum_{i=1}^{L} f(x';\theta_i)p_i^{n+1}(x)$ is constant on all possible measurement outcomes $\hat{y}$ on x. This is further equivalent to the statement that the $p_i^{n+1}(x)$'s are constants. According to the updating rule given by Equation 3.2, this is equivalent to Condition 2 in Lemma 4.2.

□

Lemma 4.3. For $\forall n \geq 0, \forall x \in \mathcal{X}$, the KGDP-$H$ score $\nu^{KGDP-H,n}(x) \geq 0$. Equality holds if and only if $p_i^n = 0$ if $f(x;\theta_i) \neq f(x;\theta^*)$.

Proof of Lemma 4.3. Suppose at time n, the probabilities of our prior candidates are $(p_1^n, p_2^n, ..., p_L^n)$. Since $x \log x$ is a convex function, by Jensen's inequality,

$$\begin{aligned}
\nu^{KGDP-H,n}(x) &= \mathbb{E}^n[\sum_{i=1}^{L} p_i^{n+1}(x)\log p_i^{n+1}(x)|S^n = s, x^n = x] - \sum_{i}^{L} p_i^n \log p_i^n \\
&= \sum_{i=1}^{L} \mathbb{E}^n[p_i^{n+1}(x)\log p_i^{n+1}(x)|S^n = s, x^n = x] - \sum_{i}^{L} p_i^n \log p_i^n \\
&\geq \sum_{i=1}^{L} \mathbb{E}^n\left[p_i^{n+1}(x)\right]\log \mathbb{E}^n\left[p_i^{n+1}(x)\right] - \sum_{i}^{L} p_i^n \log p_i^n \\
&= 0.
\end{aligned}$$

Since $x \log x$ is nonlinear, the only case for equality is that $p_i^{n+1}(x)$ is a constant, which means all functions with positive probability have the same value at x.

□

Lemma 4.4. For $\forall \omega \in \Omega_0$, $\forall x \in \mathcal{X}'(\omega)$, we have $\lim_{n\to\infty} \nu^{KGDP-f,n}(x)(\omega) = 0$, and
$\lim_{n\to\infty} \nu^{KGDP-H,n}(x)(\omega) = 0$.

Proof of Lemma 4.4. First, note that for fixed x and n, $\nu^{KGDP-f,n}(x)$ and $\nu^{KGDP-H,n}(x)$ are functions of $\vec{p} = (p_1^n, ..., p_L^n)$. In order to concentrate on $\vec{p}$, we pick any fixed $\omega \in \Omega_0$, any fixed $x \in \mathcal{X}'(\omega)$, and then denote $\nu^{KGDP-f,n}(x)$ and $\nu^{KGDP-H,n}(x)$ as $\nu_f(\vec{p})$ and $\nu_H(\vec{p})$. For this fixed x, we assume there are h ($1 \leq h \leq L$) functions equal to $f(x;\theta^*)$ at x. Without loss of generality, we assume they are $\theta_1, ..., \theta_h$. That is, $f(x;\theta_i) = f(x;\theta^*)$, for $\forall i \in [h]$.

Let $\vec{p_0} = (0, 0, ..., 0, p_{h+1}, ..., p_L)$, where $p_{h+1} + ... + p_L = 1$. By Lemma 4.2 and Lemma 4.3, $\vec{p_0}$ meets the conditions for $\nu_f(\vec{p})$ and $\nu_H(\vec{p})$ to be 0. That is, $\nu_f(\vec{p_0}) = 0$, $\nu_H(\vec{p_0}) = 0$.

Obviously, $\nu_f(\vec{p})$ and $\nu_H(\vec{p})$ are continuous in $\vec{p}$. We take $\nu_f(\vec{p})$ as an example.

$\sum_{i=1}^{L} f_i(x')p_i^n \exp\left[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}\right]$ is continuous. Therefore, we know that

$\max_{x'}\left(\sum_{i=1}^{L} f_i(x')p_i^n \exp\left[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}\right]\right)$ is also continuous. Hence,

$\int_{-\infty}^{+\infty} \max_{x'}\left(\sum_{i=1}^{L} f_i(x')p_i^n \exp\left[-\frac{(\hat{y}-f_i(x))^2}{2\sigma^2}\right]\right) d\hat{y}$ is continuous, too. Then by Equation 3.6, $\nu_f(\vec{p})$ is continuous.

Since $\vec{p}$ is defined on a compact set, by the Heine-Cantor theorem, both $\nu_f(\vec{p})$ and $\nu_H(\vec{p})$ are uniformly continuous. That is, for $\forall\epsilon > 0$, there exists $\delta > 0$, such that $\forall\vec{p_1} \neq \vec{p2}$, $|\vec{p_1} - \vec{p_2}| < \delta$, $|\nu_f(\vec{p_1}) - \nu_f(\vec{p_2})| < \epsilon$, $|\nu_H(\vec{p_1}) - \nu_H(\vec{p_2})| < \epsilon$.

For this $\omega \in \Omega_0$ and $x \in \mathcal{X}'(\omega)$, since $\lim_{n\to\infty} p_i^n(\omega) = 0$, $1 \leq i \leq h$, there exists N such that for $\forall n > N$, $p_i^n(\omega) < \delta/(L+1)$. Let $\vec{p'}_n = (0, ..., 0, p_{h+1}^n, ..., p_{L-1}^n, p_L^n + \sum_{i=1}^{h} p_i^n)$. Then $v_f(\vec{p'}) = v_H(\vec{p'}) = 0$, and $|\vec{p}_n - \vec{p'}_n| < \sqrt{h+h^2} \cdot \frac{\delta}{L+1} < \delta$.

Therefore, $\nu_f(\vec{p_n}) < \epsilon$, $\nu_H(\vec{p_n}) < \epsilon$, $\forall n > N$. Hence, for $\forall\omega \in \Omega_0$, and $x \in \mathcal{X}'(\omega)$, $\lim_{n\to\infty} \nu^{KGDP-f,n}(x) = 0$, $\lim_{n\to\infty} \nu^{KGDP-H,n}(x) = 0$.

□

Lemma 4.6. For any $\omega \in \Omega_0$, the alternatives measured infinitely often under the KGDP-f or KGDP-H policy constitute a sufficient set.

Proof of Lemma 4.6. Assume the contrary. For any $\omega \in \Omega_0$, if $\mathcal{X}_\infty(\omega)$ is not a sufficient set, then there exist other θ's that fit the true values of $\mathcal{X}_\infty(\omega)$. Assume there are h such θ's, $1 \leq h \leq L$. WLOG, assume they are θ_1,...,θ_h. So for any $l \in [h]$,

$$f(x;\theta_l) = f(x;\theta^*), \text{ for } \forall x \in \mathcal{X}_\infty(\omega).$$

By Lemma 4.1, for any $h+1 \leq l \leq L$, $\lim_{n\to\infty} p_l^n(\omega) = 0$.

We now show that for any $1 \leq l \leq h$, $\lim_{n\to\infty} p_l^n(\omega) > 0$ exists.

For this ω, let $\mathcal{L}_1^n, \mathcal{L}_2^n, ..., \mathcal{L}_L^n$ be the likelihood of the L θ's according to the first n measurements. Hence,

$$\mathcal{L}_l^n = \prod_{i=1}^{n} \exp(-\frac{(\hat{y}_x^i - f(x;\theta_l))^2}{2\sigma^2}).$$

Let T be the last time that we measure any x out of $\mathcal{X}_\infty(\omega)$. After time T, for any $l \in \{1, 2, ..., h\}$,

$$\begin{aligned}
p_l^n(\omega) &= \frac{\mathcal{L}_l^n}{\sum_{m=1}^{h}\mathcal{L}_m^n + \sum_{m=1}^{h}\mathcal{L}_m^n} \\
&= \frac{\mathcal{L}_l^T \prod_{i=T+1}^{n} \exp(-\frac{(\hat{y}_x^i - f(x;\theta^*))^2}{2\sigma^2})}{\sum_{m=1}^{h}\mathcal{L}_m^T \prod_{i=T+1}^{n} \exp(-\frac{(\hat{y}_x^i - f(x;\theta^*))^2}{2\sigma^2}) + \sum_{m=h+1}^{L}\mathcal{L}_m^T \prod_{i=T+1}^{n} \exp(-\frac{(\hat{y}_x^i - f(x;\theta_m))^2}{2\sigma^2})} \\
&= \frac{\mathcal{L}_l^T}{\sum_{m=1}^{h}\mathcal{L}_m^T + \sum_{m=h+1}^{L}\mathcal{L}_m^T \prod_{i=T+1}^{n} \frac{\exp(-\frac{(\hat{y}_x^i - f(x;\theta_m))^2}{2\sigma^2})}{\exp(-\frac{(\hat{y}_x^i - f(x;\theta^*))^2}{2\sigma^2})}}.
\end{aligned}$$

In Proof of Lemma 4.1, we have shown that $\lim_{n\to\infty}\prod_{i=T+1}^{n}\frac{\exp(-\frac{(\hat{y}_x^i-f(x;\theta_m))^2}{2\sigma^2})}{\exp(-\frac{(\hat{y}_x^i-f(x;\theta^*))^2}{2\sigma^2})}=0$. Therefore,

$$\lim_{n\to\infty}p_l^n(\omega)=\frac{\mathcal{L}_l^T}{\sum_{m=1}^{h}\mathcal{L}_m^T}>0 \text{ exists.}$$

By Lemma 4.4, for $\forall x\in\mathcal{X}_\infty(\omega)$, $\lim_{n\to\infty}\nu^{KGDP-f,n}(x)(\omega)=0$, $\lim_{n\to\infty}\nu^{KGDP-H,n}(x)(\omega)=0$. For $\forall x\in\mathcal{X}_\infty(\omega)^c$, $\lim_{n\to\infty}\nu^{KGDP-f,n}(x)(\omega)>0$, $\lim_{n\to\infty}\nu^{KGDP-H,n}(x)(\omega)>0$. Since the KGDP-$f$ (or KGDP-H) policy always chooses the x with the largest KGDP-f (or KGDP-H) score, it will measure some $x\in\mathcal{X}'(\omega)^c$ after T, which is contradictory to our assumption.

□

Theorem 4.7. Non-resampling KGDP-f with truth from prior is asymptotically optimal in finding both the optimal alternative and the correct parameter. The same holds for KGDP-H.

Proof of Theorem 4.7. Under either KGDP-f or KGDP-H policy, for any $\theta_l\neq\theta^*$, Lemma 4.6 implies that for any $\omega\in\Omega_0$, there exists $x\in\mathcal{X}_\infty(\omega)$ such that $f(x;\theta_l)\neq f(x;\theta^*)$. By Lemma 4.1, we have $\mathbb{P}(\lim_{n\to\infty}p_l^n=0)=1$. That is, $\mathbb{P}(\lim_{n\to\infty}p^n(\theta^*)=1)=1$, and $\mathbb{P}(\bar{f}(x)=f(x;\theta^*))=1$.

□

Lemma 5.1. For any $\omega\in\Omega_1$, the alternatives measured infinitely often under the KGDP-f or KGDP-H policy constitute a sufficient set. We denote this set as $\mathcal{X}_\infty(\omega)$.

Proof of Lemma 5.1. For any $\omega\in\Omega_1$, let T be the last time that we measure any x out of $\mathcal{X}_\infty(\omega)$.

Assume the contrary (thus $T>0$). Since $\mathcal{X}_\infty(\omega)$ is not a sufficient set, there exists at least one $\theta\in\mathbb{K}$, $\theta\neq\theta^*$ such that $f(x;\theta)=f(x;\theta^*)$ for $\forall x\in\mathcal{X}_\infty(\omega)$. Denote the set of these θ's as $\mathbb{K}'$.

Let $\mathcal{L}_1^n,\mathcal{L}_2^n,...,\mathcal{L}_K^n$ be the likelihood of the K θ's according to the first n measurements. If we regard $\mathbb{K}$ as the "candidate set" and let $w^n(\theta)$ denote its "probability", then as Proof of Lemma 4.6 shows, for $\forall\theta\in\mathbb{K}'\bigcup\{\theta^*\}$, $\lim_{n\to\infty}w^n(\theta)=\frac{\mathcal{L}_\theta^T}{\sum_{i=1}^{K}\mathcal{L}_i^T}>0$. For all other θ's, by Lemma 4.1, $\lim_{n\to\infty}w^n(\theta)=0$. $w^n(\theta)$ is actually the weight of θ when we do resampling. Hence, as n gets larger, θ's in $\mathbb{K}'\bigcup\{\theta^*\}$ always rank higher than others to enter the small pool (i.e, the sub-level set), from which we resample.

We refer to the event as A_n that at least two θ's in $\mathbb{K}'\bigcup\{\theta^*\}$ are in the candidate set at time n. We first show that A_n happens infinitely often. That is, $\mathbb{P}_\theta(\bigcap_{m\geq1}\bigcup_{n\geq m}A_n)=1$, where $\mathbb{P}_\theta$ indicates the fact that this probability is calculated in the $(\Omega_\theta,\mathcal{F}_\theta,\mathbb{P}_\theta)$ space.

There are two cases at a resampling step: (1) $|\mathbb{L}_{rm}|\geq1$ or (2) $|\mathbb{L}_{rm}|=1$ (remember that L_{rm} is the set of candidates to remove). Since resampling happens infinite often, at least either (1) or (2) happens infinitely often.

We first assume (1) happens infinitely often, indexed by a subsequence of times $\{s_n\}$. Let B_{s_n} be the event that at time s_n, two θ's from $\mathbb{K}'\bigcup\{\theta^*\}$ are selected in

the first draw. (Remember that we may have multiple draws if resampling keeps being triggered. However, for n large enough, if two such θ's are selected in the first draw, they will not be dropped in this iteration. Hence, $B_{s_n} \subset A_{s_n}$) Since all $w^n(\theta)$'s have limits, we know $\lim_{n\to\infty} \mathbb{P}_\theta(B_{s_n}) > 0$. Therefore, $\sum_{n=1}^{\infty} \mathbb{P}_\theta(B_{s_n}) = \infty$. On the other hand, for a fixed ω, $w^n(\theta)$'s are all decided, and therefore B_{s_n} are pairwise independent. By the Borel-Cantelli Lemma, the probability that B_{s_n} happens infinitely often, i.e, $\mathbb{P}_\theta(\bigcap_{m\geq 1} \bigcup_{s_n \geq m} B_{s_n}) = 1$. We know that B_{s_n} is a subset of A_{s_n}. Hence, $\mathbb{P}_\theta(\bigcap_{m\geq 1} \bigcup_{s_n \geq m} A_{s_n}) = 1$, and furthermore, $\mathbb{P}_\theta(\bigcap_{m\geq 1} \bigcup_{n \geq m} A_n) = 1$.

Otherwise, if (1) happens a finite number of times and (2) happens infinitely often indexed by $\{t_n\}$, we assume the contrary. That is, there exists time T_2, such that after T_2, among the L candidates there is at most one $\theta \in \mathbb{K}' \bigcup \{\theta^*\}$. At time t_n, since $|\mathbb{L}_{rm}| = 1$, by assumption, the other $L-1$ candidates must be identical, and the newly drawn one must be the same as the other $L-1$ ones. The probability that this keeps happening from time t_M to infinity, is $\lim_{N\to\infty} \prod_{n=M}^{N} w^{t_n}(\theta) = 0$. Therefore, our assumption does not hold and A_n stills happens infinitely often!

Hence, for any fixed $\omega \in \Omega_1$, with probability 1, we have the subsequence $\{q_n(\omega)\}$, such that at these times, the candidates contain at least two θ's in $\mathbb{K}' \bigcup \{\theta^*\}$, denoted as θ_1^n and θ_2^n. For any $x \in \mathcal{X}_\infty(\omega)$, $\lim_{n\to\infty} \nu^{KGDP-f,n}(x)(\omega) = 0$, $\lim_{n\to\infty} \nu^{KGDP-H,n}(x)(\omega) = 0$; while for any x such that $f(x;\theta_1^n) \neq f(x;\theta_2^n)$ and any q_n, $\nu^{KGDP-f,q_n}(x)(\omega) > 0$, $\nu^{KGDP-H,q_n}(x)(\omega) > 0$. As there are only finite combinations of θ_1 and θ_2 and each $\theta \in \mathbb{K}' \bigcap \{\theta^*\}$ has a positive probability in the limit, there exists $\epsilon > 0$, such that $\sup \nu^{KGDP-f,q_n}(x)(\omega) > \epsilon$, $\nu^{KGDP-H,q_n}(x)(\omega) > \epsilon$. Since q_n can be larger than T, this is contradictory to the fact that KGDP-f or KGDP-H will not measure any x out of $\mathcal{X}_\infty(\omega)$ after time T.

□

Theorem 5.2. Resampling KGDP-f is asymptotically optimal in finding the optimal alternative and the correct parameter. The same holds for KGDP-H.

Proof of Theorem 5.2. For any fixed $\omega \in \Omega_1$, let C_n ($C_n \in \mathcal{F}_\theta$) be the set of events that the L candidates includes θ^* at time n.

By Lemma 5.1, θ^* is the only parameter in $\mathbb{K}$ that fits $\mathcal{X}_\infty(\omega)$. Assume $\{s_n\}$ is the subsequence of times when resampling happens. Then $\mathbb{P}_\theta(C_{s_n}) > \frac{w^{s_n}(\theta^*)(\omega)}{\sum_{i=1}^{K} w_i^{s_n}(\omega)}$, where the right hand side is the probability that θ^* is chosen in the first draw.

We show $\mathbb{P}_\theta(\bigcup_{m\geq 1} \bigcap_{n\geq m} C_n) = 1$, which means that θ^* is in the candidate set for all but a finite number of times.

Since resampling only happens on $\{s_n\}$, we know $\mathbb{P}_\theta(\bigcup_{m\geq 1} \bigcap_{n\geq m} C_n) = \mathbb{P}_\theta(\bigcup_{m\geq 1} \bigcap_{n\geq m} C_{s_n})$. Intuitively, this means if after time T, every resampling step chooses θ^* for the candidate set, then θ^* appears in the candidate set every iteration after T.

For this fixed ω, we know

$$\mathbb{P}_\theta(C_{s_n}^c) = 1 - \mathbb{P}_\theta(C_{s_n}) < \frac{\sum_{\theta\in\mathbb{K}, \theta\neq\theta^*} w^{s_n}(\theta)}{\sum_{\theta\in\mathbb{K}} w^{s_n}(\theta)} < \frac{\sum_{\theta\in\mathbb{K}, \theta\neq\theta^*} w^{s_n}(\theta)}{w^{s_n}(\theta^*)}.$$

From proof of Lemma 4.1 (for detail, check out Equation (A.1) and (A.2)), for this $\omega \in \Omega_1$, we know there exists N_1, such that for all $n > N_1$ and any $\theta \neq \theta^*$,

$$\frac{w^{s_n}(\theta)}{w^{s_n}(\theta^*)} < \exp(-n\frac{(f^* - f_l)^2}{4\sigma^2}).$$

Assume $\sup_{l,x} \frac{(f^*-f_l)^2}{4\sigma^2} = \rho$. Hence, for $n > N_1$, $\mathbb{P}_\theta(C_{s_n}^c) < Ke^{-n\rho}$. Therefore,

$$\sum_{n\geq 1} \mathbb{P}_\theta(C_{s_n}^c) < \sum_{n\geq 1} \mathbb{P}_\theta(D_n^c) < N_1 + \sum_{n\geq N_1+1} Ke^{-n\rho} = N_1 + K\frac{e^{-(N_1+1)\rho}}{(1-e^{-\rho})} < \infty.$$

By the Borel-Cantelli Lemma, $\mathbb{P}_\theta(\bigcap_{m\geq 1}\bigcup_{n\geq m} C_{s_n}^c) = 0$. Hence, $\mathbb{P}_\theta(\bigcup_{m\geq 1}\bigcap_{n\geq m} C_n) = \mathbb{P}_\theta(\bigcup_{m\geq 1}\bigcap_{n\geq m} C_{s_n}) = 1$.

Hence, for any $\omega \in \Omega_1$, with probability 1, there exists $T(\omega)$, such that for any $n \geq T(\omega)$, θ^* is in the candidate set, and $\lim_{n\to\infty} p^n(\theta^*)(\omega) = 1$.

This holds for every $\omega \in \Omega_1$, hence, $\mathbb{P}_{\Omega\times\Omega_\theta}\{\lim_{n\to\infty} p^n(\theta^*) = 1\} = 1$, where $\mathbb{P}_{\Omega\times\Omega_\theta}$ indicates the probability in the full probability space $(\Omega_f, \mathcal{F}_f, \mathbb{P}_f)$, which is the product space of $(\Omega, \mathcal{F}, \mathbb{P})$ and $(\Omega_\theta, \mathcal{F}_\theta, \mathbb{P}_\theta)$. Since we find θ^* with probability 1, we also find the optimal x with probability 1.

□

REFERENCES

[1] J.-Y. Audibert and S. Bubeck, *Best Arm Identification in Multi-Armed Bandits*, in COLT 2010 - Proc., 2010, p. 13 p.

[2] P. Auer, N. Cesa-Bianchi, and P. Fischer, *Finite-time analysis of the multiarmed bandit problem*, Mach. Learn., 47 (2002), pp. 235–256, doi:10.1023/A:1013689704352.

[3] R. R. Barton and M. Meckesheimer, *Chapter 18 Metamodel-Based Simulation Optimization*, Handbooks Oper. Res. Manag. Sci., 13 (2006), pp. 535–574, doi:10.1016/S0927-0507(06)13018-2.

[4] L. Bianchi, M. Dorigo, L. M. Gambardella, and W. J. Gutjahr, *A survey on metaheuristics for stochastic combinatorial optimization*, Nat. Comput., 8 (2009), pp. 239–287, doi:10.1007/s11047-008-9098-4.

[5] C. H. Chen, J. Lin, E. Yücesan, and S. E. Chick, *Simulation budget allocation for further enhancing the efficiency of ordinal optimization*, Discret. Event Dyn. Syst. Theory Appl., 10 (2000), pp. 251–270, doi:10.1023/A:1008349927281.

[6] C.-H. C. C.-H. Chen, *An effective approach to smartly allocate computing budget for discrete event simulation*, Proc. 1995 34th IEEE Conf. Decis. Control, 3 (1995), pp. 2598–2603, doi:10.1109/CDC.1995.478499.

[7] S. Chen, K.-r. G. Reyes, M. Gupta, M. C. Mcalpine, and W. B. Powell, *Optimal learning in experimental design using the knowledge gradient policy with application to characterizing nanoemulsion stability*, SIAM/ASA J. Uncertain. Quantif., 3 (2014), pp. 320–345, doi:10.1137/140971129.

[8] S. E. Chick, *Subjective Probability and Bayesian Methodology*, Handbooks Oper. Res. Manag. Sci., 13 (2006), pp. 225–257, doi:10.1016/S0927-0507(06)13009-1.

[9] S. E. Chick and K. Inoue, *New Two-Stage and Sequential Procedures for Selecting the Best Simulated System*, Oper. Res., 49 (2001), pp. 732–743, doi:10.1287/opre.49.5.732.10615.

[10] P. I. Frazier, W. B. Powell, and S. Dayanik, *A Knowledge-Gradient Policy for Sequential Information Collection*, SIAM J. Control Optim., 47 (2008), pp. 2410–2439, doi:10.1137/070693424.

[11] P. I. Frazier, W. B. Powell, and S. Dayanik, *The knowledge-gradient policy for correlated normal beliefs*, INFORMS J. Comput., 21 (2009), pp. 599–613, doi:10.1287/ijoc.1080.0314.

[12] M. C. Fu, *Chapter 19 Gradient Estimation*, Handbooks Oper. Res. Manag. Sci., 13 (2006), pp. 575–616, doi:10.1016/S0927-0507(06)13019-4.

[13] J. Gittins, K. Glazebrook, and R. Weber, *Multi-armed Bandit Allocation Indices*, 2011, doi:10.1002/9780470980033.

[14] R. B. Gramacy and H. K. H. Lee, *Optimization under unknown constraints*, Bayesian Stat. 9, (2011), p. 229.

[15] S. S. Gupta and K. J. Miescke, *Bayesian look ahead one-stage sampling allocations for selection of the best population*, J. Stat. Plan. Inference, 54 (1996), pp. 229–244.

[16] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, no. 2, 2009, doi:10.1007/b94608.

[17] D. He, S. E. Chick, and C.-H. Chen, *Opportunity Cost and OCBA Selection Procedures in Ordinal Optimization for a Fixed Number of Alternative Systems*, IEEE Trans. Syst. Man Cybern. Part C, Appl. Rev., 37 (2007), pp. 951–961, doi:10.1109/TSMCC.2007.900656.

[18] L. Hong and B. L. Nelson, *A brief introduction to optimization via simulation*, Winter Simul. Conf., (2009), pp. 75–85.

[19] H. Jiang and U. V. Shanbhag, *On the solution of stochastic optimization and variational problems in imperfect information regimes*, ArXiv e-prints, (2014), arXiv:1402.1457.

[20] D. Jones, M. Schonlau, and W. Welch, *Efficient global optimization of expensive black-box functions*, J. Glob. Optim., 13 (1998), pp. 455–492, doi:10.1023/a:1008306431147.

[21] L. P. Kaelbling, *Learning in Embedded Systems*, MIT Press, Cambridge, MA, 1993.

[22] S.-H. Kim and B. L. Nelson, *Selecting the best system*, vol. 13 of Handbooks in Operations Research and Management Science, Elsevier, 2006, doi:10.1016/S0927-0507(06)13017-0.

[23] D. M. Negoescu, P. I. Frazier, and W. B. Powell, *The knowledge-gradient algorithm for sequencing experiments in drug discovery*, INFORMS J. Comput., 23 (2011), pp. 346–363, doi:10.1287/ijoc.1100.0417.

[24] S. Ólafsson, *Chapter 21 Metaheuristics*, Handbooks Oper. Res. Manag. Sci., 13 (2006), pp. 633–654, doi:10.1016/S0927-0507(06)13021-2.

[25] W. B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, John Wiley and Sons, New York, 2 ed., 2011.

[26] S. Singh, T. Jaakkola, M. L. Littman, C. Szepes, and A. S. Hu, *Convergence Results for Single-Step On-Policy Reinforcement-Learning Algorithms*, Mach. Learn., 39 (2000), pp. 287–308, doi:10.1023/A:1007678930559.

[27] J. Spall, *Introduction to Stochastic Search and Optimization*, Wiley, New York, 1 ed., 2003.

[28] M. J. Streeter and S. F. Smith, *Principles and Practice of Constraint Programming - CP 2006: 12th International Conference, CP 2006, Nantes, France, September 25-29, 2006. Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, ch. A Simple D, pp. 560–574, doi:10.1007/11889205_40, http://dx.doi.org/10.1007/11889205{_}40.

[29] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, MA, 1998.